

Checking and Distributing Statistical Model-Checking

Peter Bulychev

Alexandre David

Kim G. Larsen

Axel Legay

Marius Mikucionis

Danny Bøgsted Poulsen

Outline

- UPPAAL-SMC in a Nutshell
- Distributing SMC
- Checking DSMC
- Case-studies

UPPAAL

Safety ✓

$A[] \text{ forall } (i : \text{id_t}) \text{ forall } (j : \text{id_t})$
 $\text{Train}(i).\text{Cross} \ \&\& \ \text{Train}(j).\text{Cross} \text{ imply } i == j$

Reachability ✓

$E \leftrightarrow \text{Train}(0).\text{Cross} \text{ and } \text{Train}(1).\text{Stop}$

Liveness ✓

$\text{Train}(0).\text{Appr} \rightarrow \text{Train}(0).\text{Cross}$

$A \leftrightarrow \dots E[] \dots$ ✓

Limited quantitative analysis ✓

sup: .. inf: ..

Performance properties ✗

State-space explosion ✗

UPPAAL SMC

Performance properties

$\Pr[\leq 200](\neg \text{Train}(5).\text{Cross})$

$\Pr[\leq 100](\neg \text{Train}(0).\text{Cross}) \geq 0.8$



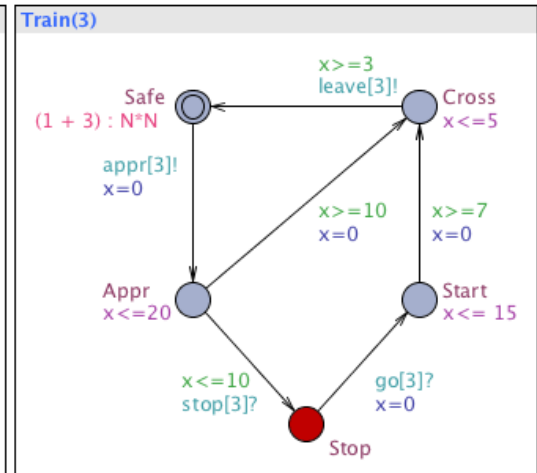

$\Pr[\leq 100](\neg \text{Train}(5).\text{Cross}) \geq$
 $\Pr[\leq 100](\neg \text{Train}(1).\text{Cross})$

State-space explosion

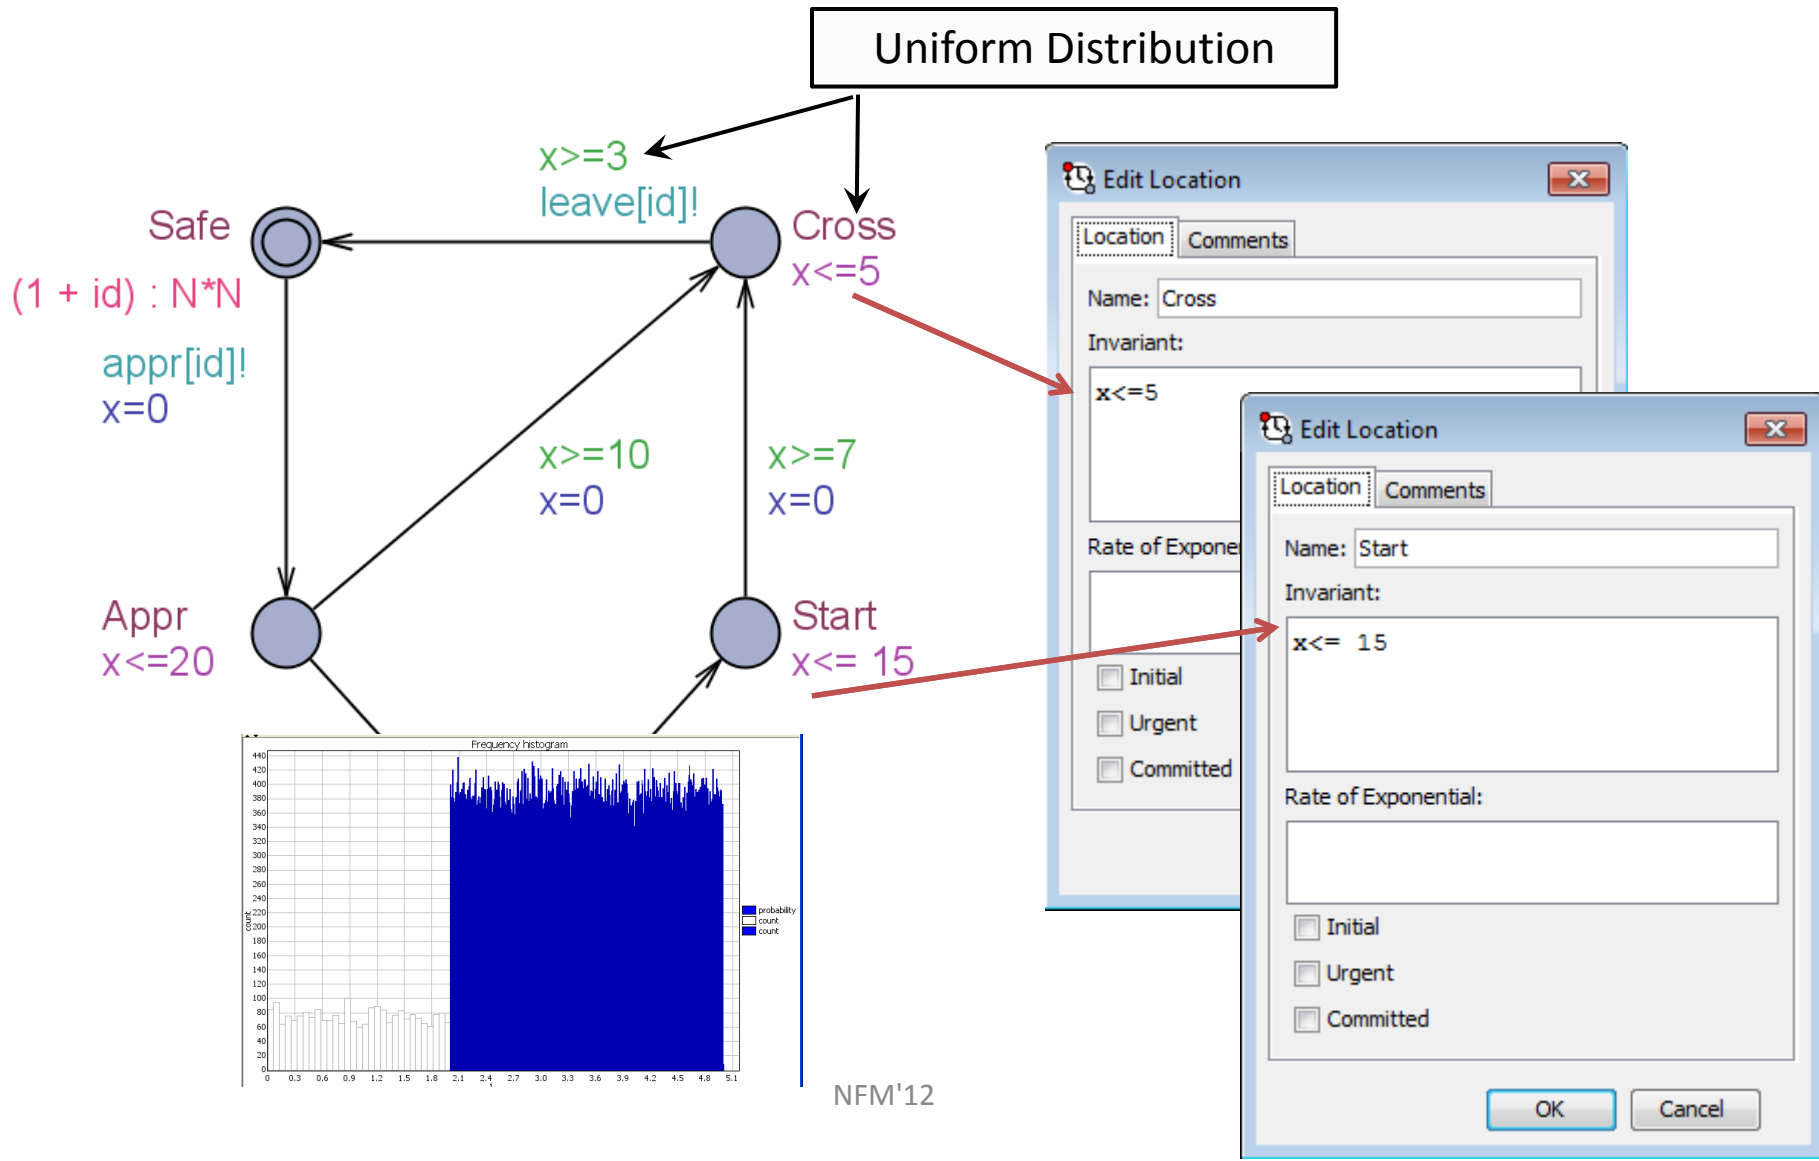
Generate random runs

Performance properties

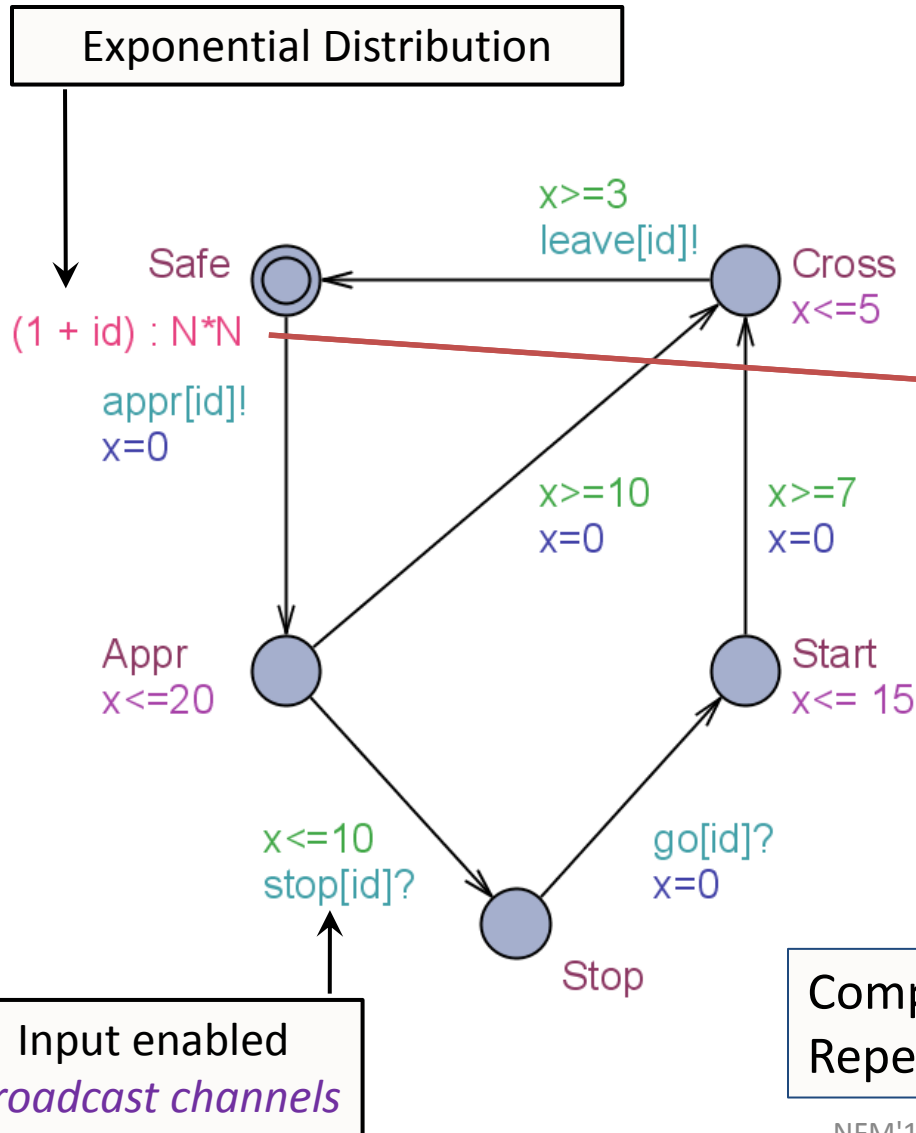
State-space explosion



Stochastic Semantics of UPPAAL TA



Stochastic Semantics of UPPAAL TA



Edit Location

Location Comments

Name: Safe


Invariant:

Rate of Exponential:

(1 + id) : N*N

☒ Initial

☐ Urgent



Composition =
Repeated races between components

Queries

Syntax

- Hypothesis testing

`Pr [<= 100] (< > expr) >= 0.1`

`x <= 100 # <= 50 [] expr <= 0.5`

- Evaluation

`Pr [<= 100] (< > expr)`

- Comparison

`Pr [<= 20] (< > e1) >= Pr [<= 10] (< > e2)`

- Expected value

`E [<= 10 ; 1000] (min : expr)`

Explicit number of runs. Min or max.

- Simulations

`simulate 10 [<= 100] { expr1 , expr2 }`

Queries

Syntax

- Hypothesis testing

`Pr[<=100] (<> expr) >=0.1`

`x<=100 #<=50 [] expr <=0.5`

- Evaluation

`Pr[<=100] (<> expr)`

- Comparison

`Pr[<=20] (<> e1) >=Pr[<=10] (<> e2)`

- Expected value

`E[<=10;1000] (min: expr)`

Explicit number of runs. Min or max.

- Simulations

`simulate 10 [<=100] {expr1,expr2}`

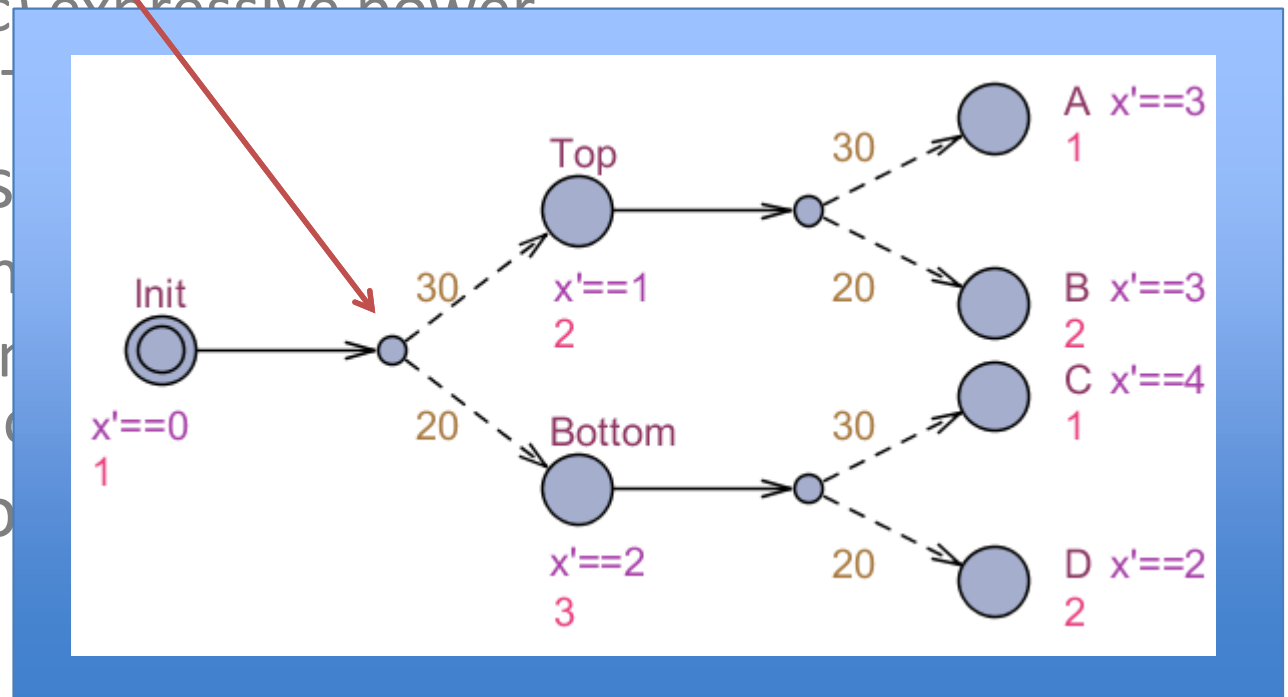
SMC in UPPAAL

Invariants:
 $x'==0 \ \&\& \ y'==\text{bool_fun}()$

- Constant Slope Timed Automata
 - **Clocks** may have different (integer) **slope** in different locations.
 - Branching edges with discrete probabilities (weights).
 - Beyond Priced TA, Energy TA. Equal LHA in (non-stochastic) expressive power.
 - Beyond DTMC, beyond CTMC (with multiple rewards)
- All features of UPPAAL supported
 - User defined functions and types
 - Expressions in guards, invariants, clock-rates, delay-rates (rationals), and weights.
- New GUI for plot-composing and exporting.



SMC in UPPAAL

- Constant Slope Timed Automata
 - Clocks may have different (integer) slope in different locations.
 - **Branching edges** with discrete probabilities (weights).
 - Beyond Priced TA, Energy TA. Equal LHA in (non-stochastic) expressive power.
 - Beyond DT
- All features
 - User defin
 - Expression
 - rates (ratio
- New GUI fo



SMC in UPPAAL

Constant Slope Timed Automata




rates (rationals), and weights.

- New GUI for plot-composing and exporting.

Distributing SMC

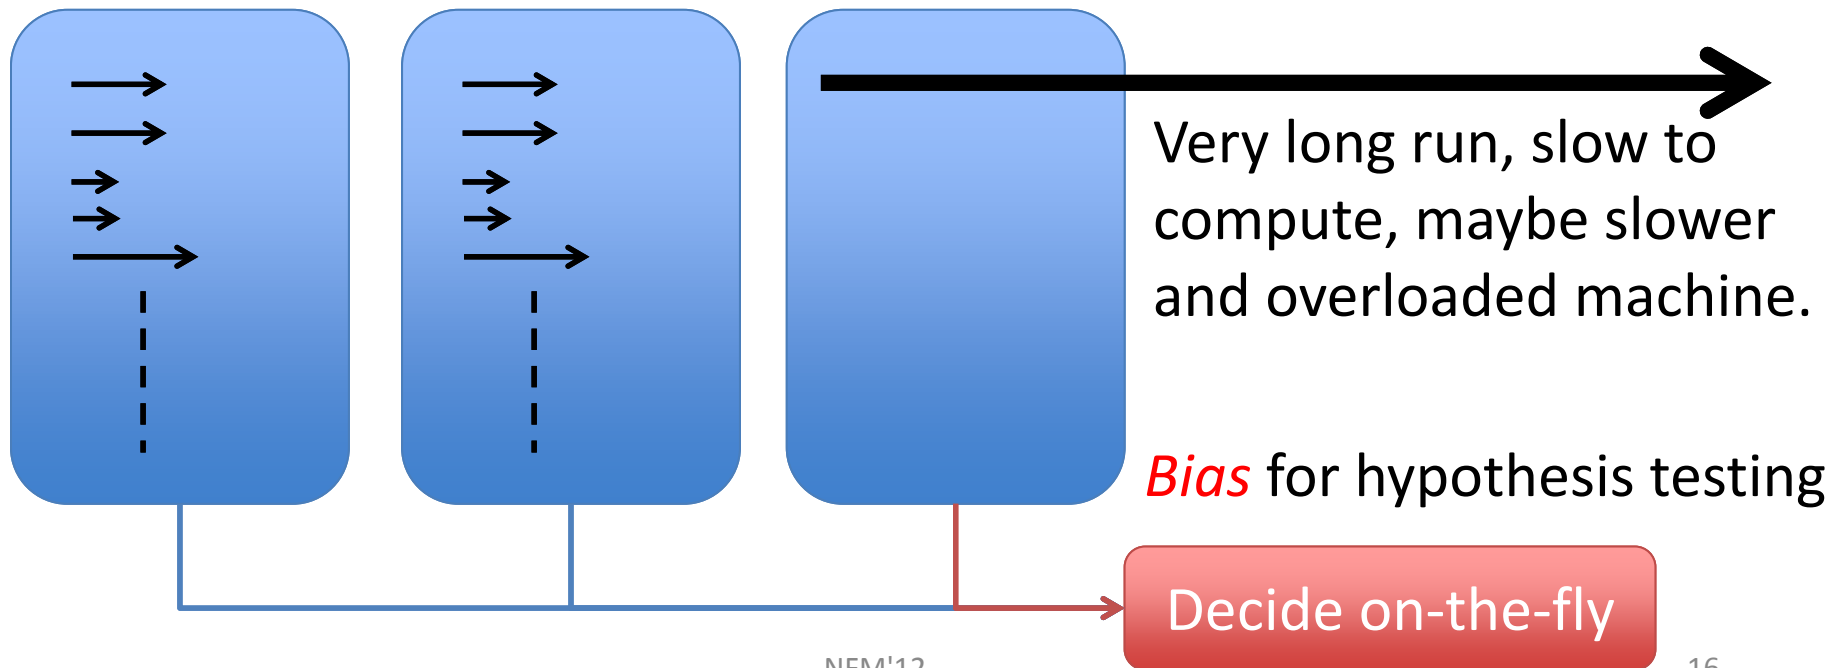
- Distributed SMC
 - Evaluation – trivial to parallelize
 - Hypothesis – careful

Hypothesis testing:
intuition.

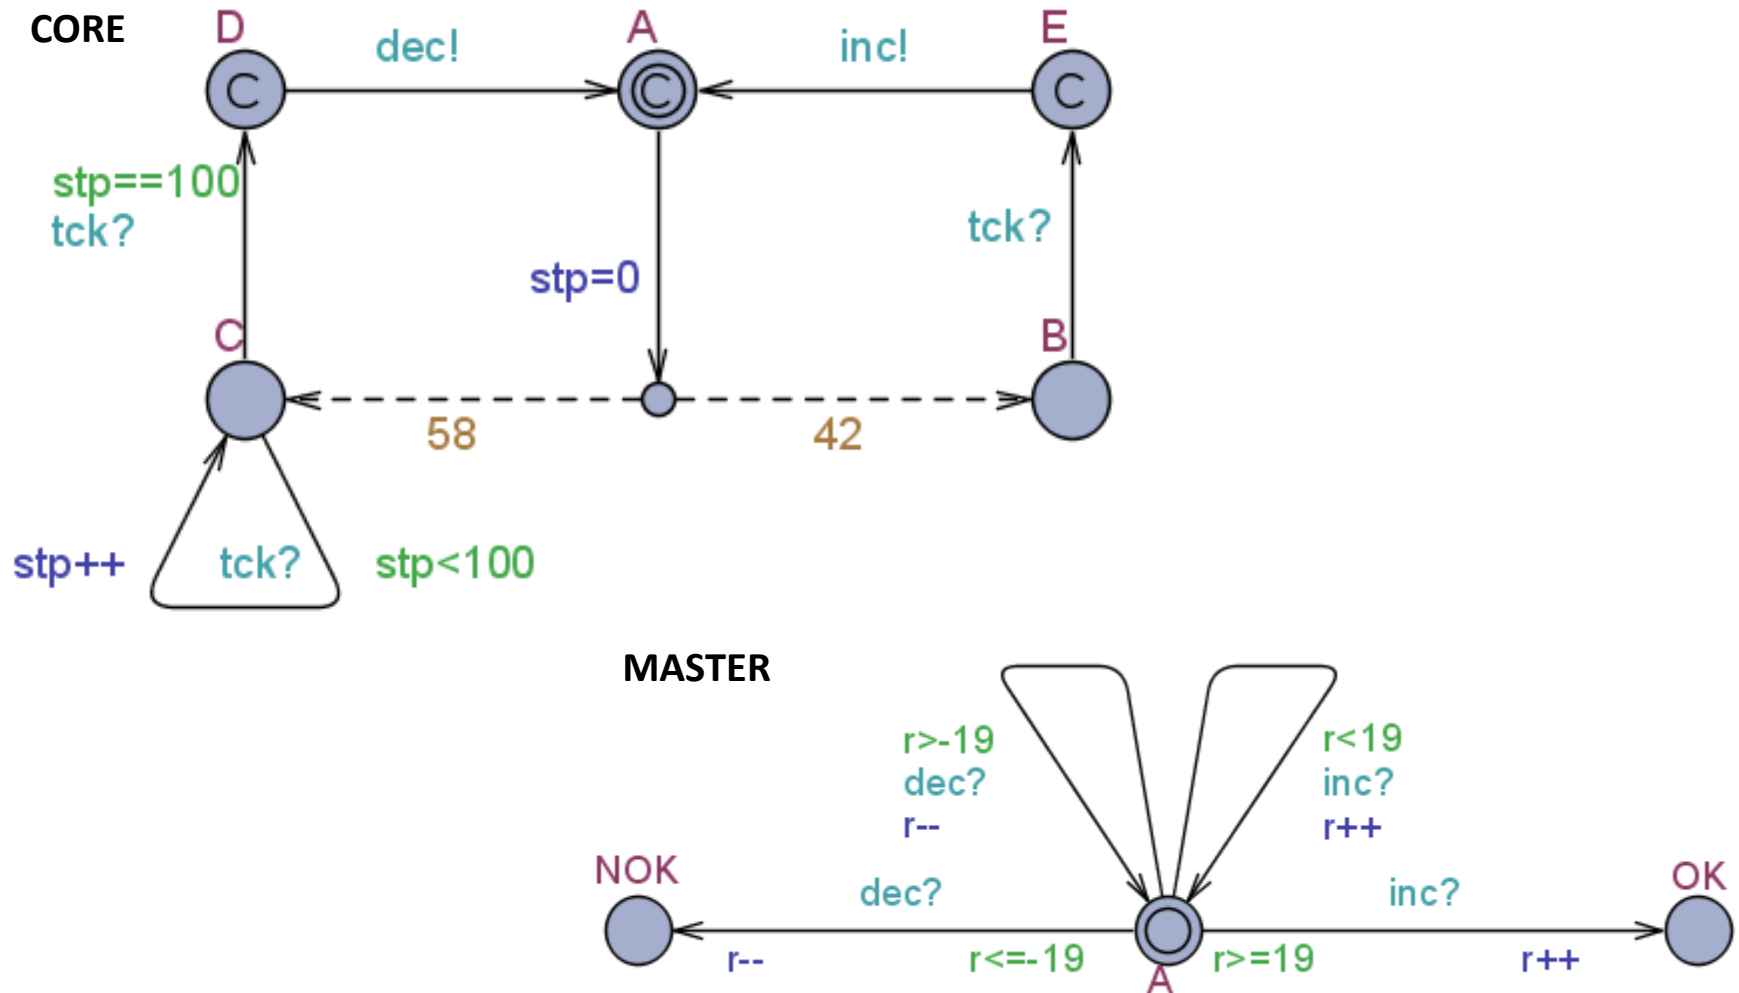


Distributing SMC

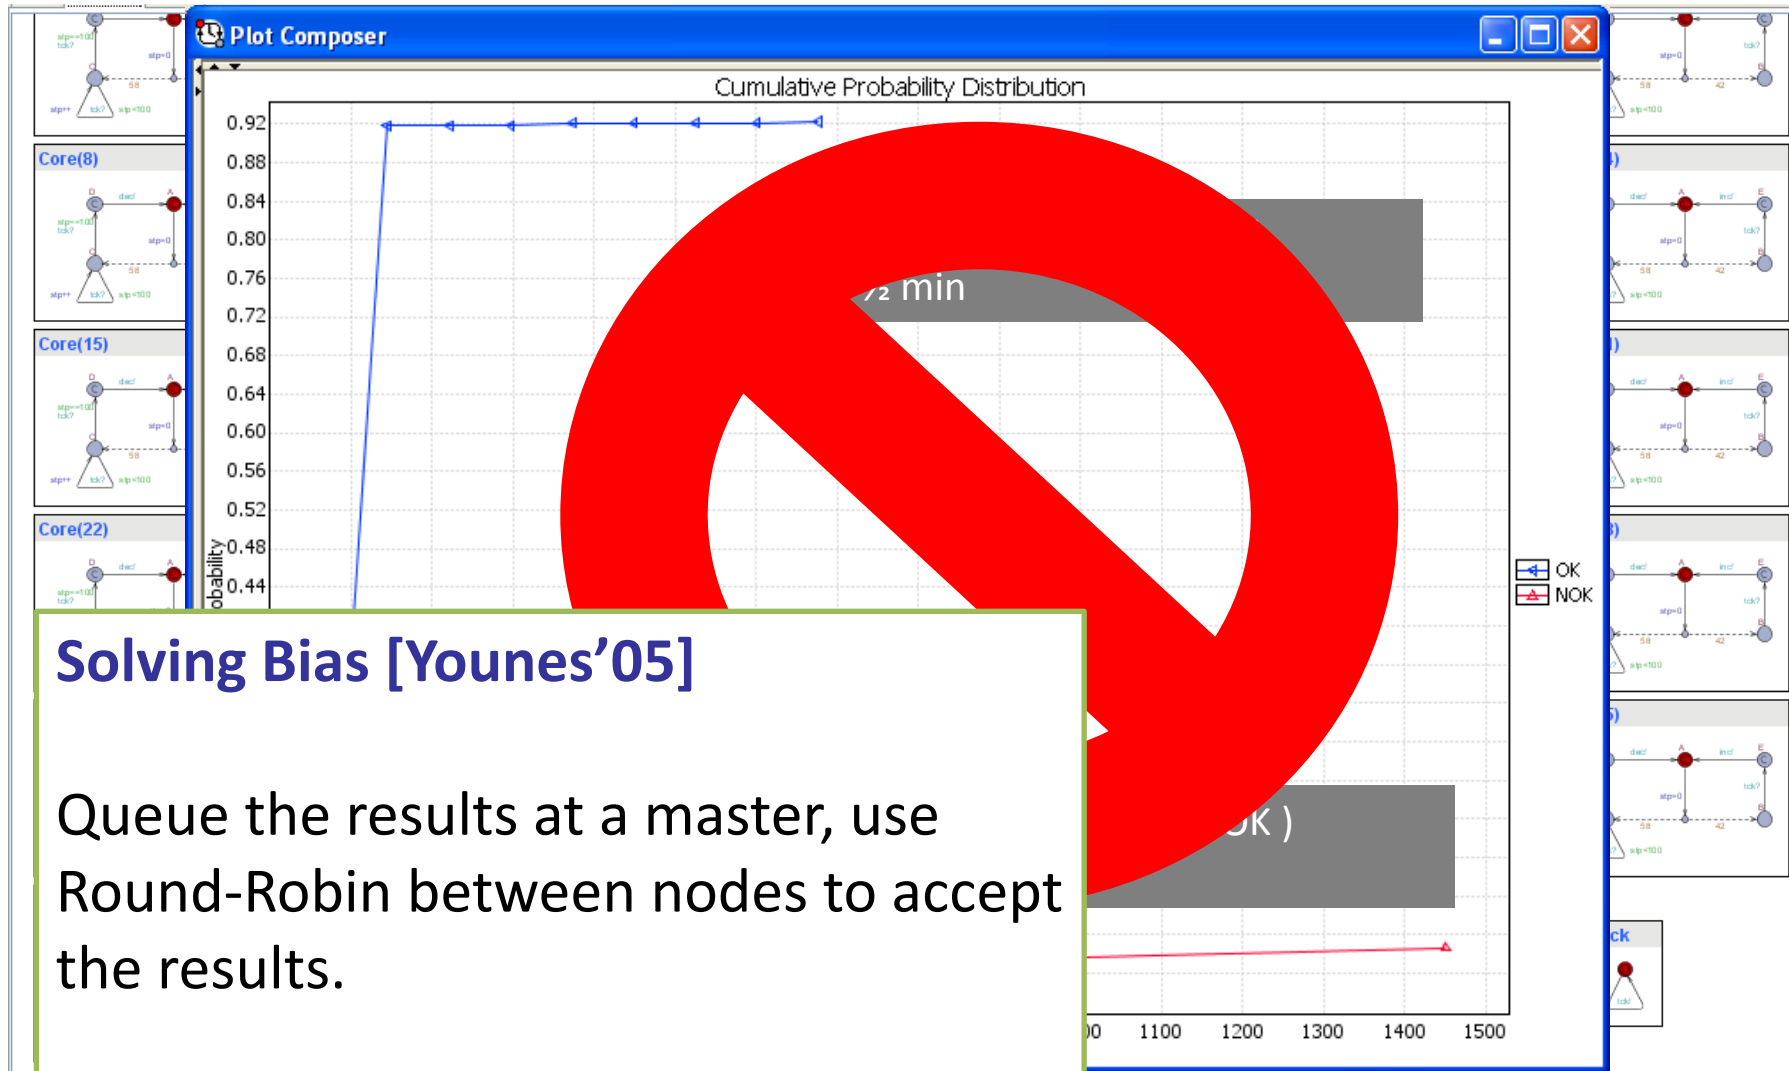
- Distributing hypothesis testing.



Distributing SMC – Naïve Approach



Distributing SMC – Naïve Approach




Solving Bias [Younes'05]

Queue the results at a master, use Round-Robin between nodes to accept the results.


Our Implementation

- Use a batch of **B** (e.g 10) runs, transmit one count per batch.
- Use asynchronous communication (MPI)
- Queue results at the master and wait only when the buffer (size=**K**) is full.



Our Implementation

- Senders have a buffer of (K) asynchronously sent messages and blocks only when the buffer is full.
- The master periodically add results in the buffer.




Results


16, 32, 128 cores, Vary Buffer & Batch Sizes

“Small” model: Exhibit expected behaviour.


16 cores




32 cores




128 cores




16 cores



32 cores



128 cores



“Big” model: Amortize the differences.


Checking DSMC

- We can model the algorithm inside UPPAAL.
 - Run SMC on it, even DSMC!

```
1 // buffer portion for early termination:
2 const int P = (K<=4)?K : ((K<=8)?5 : ((K<=16)?8 : ((K<=32)?10 : 12)));
3 bool H0 = false, H1 = false; // for hypothesis H0 and H1
4 int batch[N][K]; // buffer of batches (K batches for N nodes)
5 long satisfied =0, unsatisfied =0; // information about filled lines
6 long sat=0, unsat=0, unknown=N*P*B; // early results in unfilled lines
7 long logRatio = 0, ratioLow = 0, ratioUp = 0; // scaled by p.scale
```


Slave Node

train gate model



$Pr[\# \leq 100](\langle \rangle \text{Train}(5).\text{Cross})$

Frequency Histogram



start



$level[id] < K$
req?
 $runs=0, sat=0,$
 $x=0$

tmp



$x \geq \text{LatencyLow}$
 $deliver[id]!$
 $busy[id]=0,$
 $value=sat$

$i:bucket_t$
 $j=i, runs++,$
 $busy[id]=1,$
 $sat+=(i < H_last),$
 $x=0$

$w[i]$



compute
 $x \leq (j+1) * H_step$

$x \geq j * H_step$
 $j=0, x=0$

$runs < B$

$runs == B$

iterate

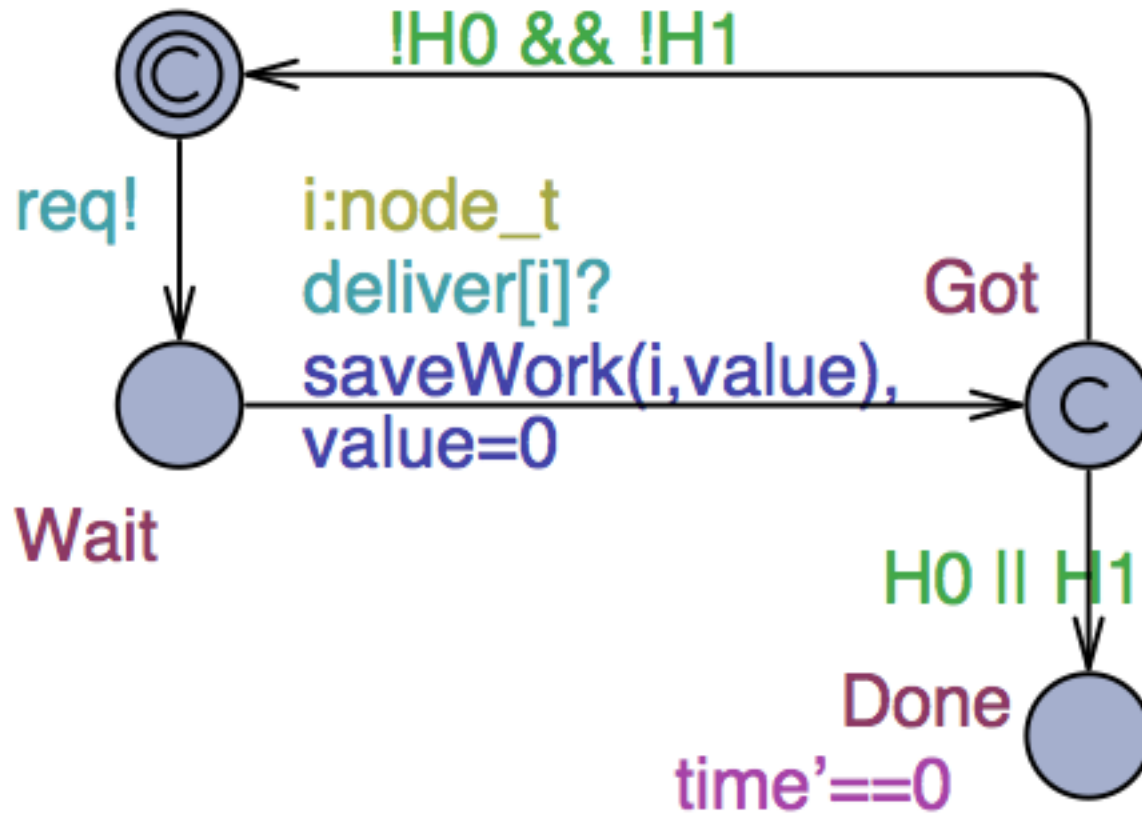


$x \leq \text{LatencyUp}$

latency



Master Node



```

8 void saveWork(const node_t node, const int value) {
9     if (level [node]<=P) { // entered the early results portion
10         sat += value; unsat += B-value; unknown -= B;
11     }
12     batch[node][level [node]] = value; level [node]++; // store
13     if (level [node]==1) { // entered at the lowest level
14         bool filled = forall (i: node.t) level [i]>0;
15         if (filled ) { // line at the lowest level has been filled
16             int L;
17             for (i: node.t) { // shift all queues one by one
18                 satisfied += batch[i][0]; // count as firm results
19                 unsatisfied += B-batch[i][0];
20                 sat -= batch[i][0]; // discount from early results
21                 unsat -= B-batch[i][0]; unknown += B;
22                 level [i]--; // remove from buffer
23                 for (L=0; L<level[i]; ++L) {
24                     batch[i][L] = batch[i][L+1]; // shift
25                     if (L==P) { // entered the early results portion
26                         sat += batch[i][L+1]; unsat += B-batch[i][L+1];
27                     }
28                 }
29                 batch[i][level [i]]=0; // cleanup
30             }
31             logRatio = p.valAcc*satisfied + unsatisfied *p.valRef;
32             if (logRatio <= p.logInf) H0 = true;
33             if (logRatio >= p.logSup) H1 = true;
34         }
35     }
36     ratioLow = p.valAcc*(satisfied +sat+unknown) +
37                 p.valRef*(unsatisfied +unsat);
38     ratioUp = p.valAcc*(satisfied +sat) +
39                p.valRef*(unsatisfied +unsat+unknown);
40     if (ratioUp <= p.logInf) H0 = true;
41     if (ratioLow >= p.logSup) H1 = true;
42 }

```


Read batch.

Exact decision.

Safe approximate decision.

Results

N=16
B=1..10
K=1,2,4,8



Can predict performance.
Can derive more information,
e.g., processor usage.
Validate implementation.

Property used:

$E[\text{time} \leq 1000; 1000]$ (max: usage)

Conclusions:


K=1 has huge effect and should be avoided.

K=2 has effect if $B < 20$.


$K > 2$ are indistinguishable on homogeneous cluster.

**$K > 2$ and $B > 20$: number of simulations scale
linearly to the number of cores used.**



Case Studies




FIREWIRE



BLUETOOTH



LMAC



LMAC

Lightweight Media Access Control (LMAC)

- Problem domain:
 - communication scheduling
- Targeted for:
 - self-configuring networks,
 - collision avoidance,
 - low power consumption
- Application domain:
 - wireless sensor networks

LMAC Protocol Design

- Four phases:
 - Initialization (listen until a neighbor is heard)
 - Waiting (delay a random amount of time frames)
 - Discovery (wait for entire frame and note used slots)
 - Active
 - choose free slot,
 - use it to transmit, including info about detected collisions
 - listen on other slots
 - fallback to Discovery if collision is detected
- Only neighbors can detect collision and tell the user-node that its slot is used by others

added power

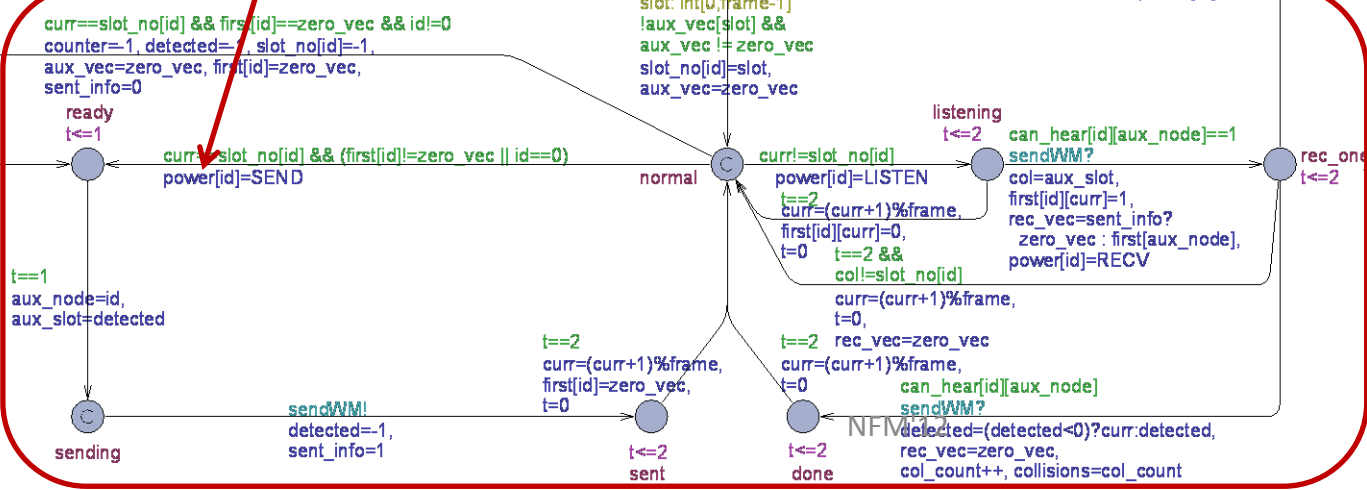
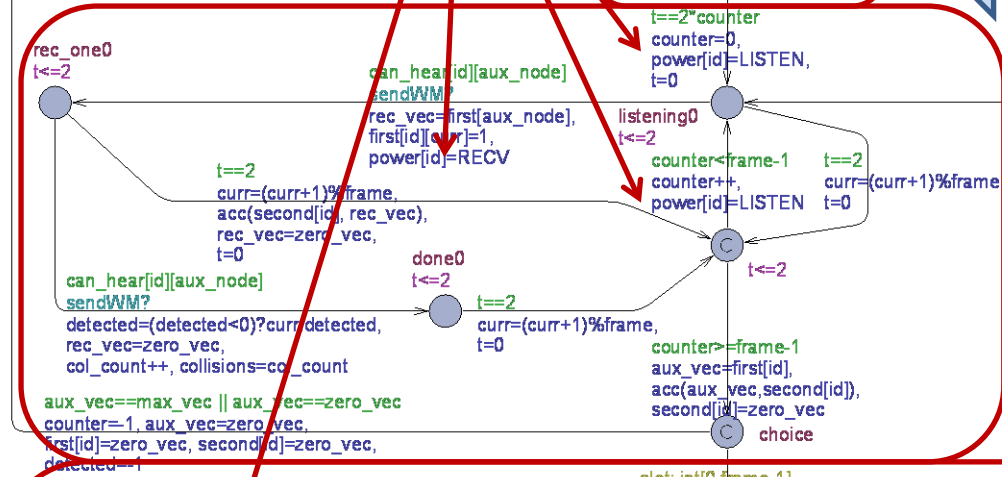

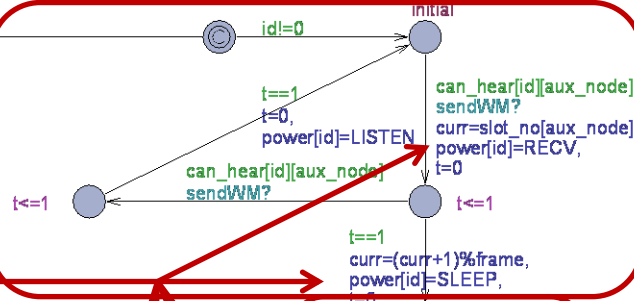
initialization

random wait

discovery

active usage

id==0
curr=0,
power[id]=SEND




Classical vs. Statistical MC

- A.Fehnker, L.v.Hoesel and A.Mader used UPPAAL to explore 4- and 5-node topologies and found cases with **perpetual collisions**.
- However they could not know whether the next collisions are inevitable.
- Statistical MC offers an insight by calculating the **probability over the number of collisions**.
+ estimated cost in terms of energy.

LMAC Simple Statistics for 4 Nodes

- Wait distribution:

- geometric



- uniform




- Network topology:

- chain






- ring



- Collision probability

- Collision count

- Power consumption

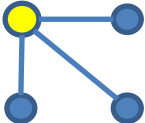
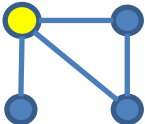
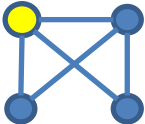
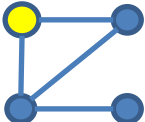


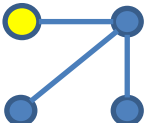
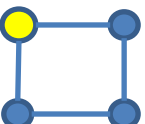

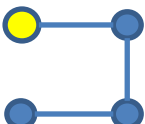
NFM'12
Pr[collisions<=50000] (<= time>=1000)

LMAC with Parameterized Topology



Collision probability in a 4 node network of a randomly generated topology:

$\Pr[\text{time} \leq 200] (< \text{col_count} > 0)$


	topology	collision probability
(star)		[0.36; 0.39]
		[0.29; 0.36]
		[0.26; 0.30]
		[0.19; 0.21]

	topology	collision probability
(ring)		[0.08; 0.19]
		[0.11; 0.13]
		[0.08; 0.15]
		[0.049; 0.050]

10-Node Chain




The first collisions can be as late as 800tu.
It is very likely (>94%) that
there will be 0 collisions.
But if they happen, some are perpetual.




10-Node Ring


Probability Density Distribution



Probability Density Distribution





Probability Density Distribution




The first collisions can be as late as 920tu.
It is very likely (>90%) that
there will be 0 collisions.
But if they happen, they are perpetual.

10-Node Star



The first collisions happen before 500tu.
It is unlikely (8.2%) that
there will be 0 collisions.
And if they happen, they are perpetual.



10-Node Random Topologies

Generated **10000** random topologies

Checked the property:

$\Pr[\text{time} \leq 2000](\neq \text{col_count} > 42)$

(perpetual collisions are likely)

One instance on a laptop takes \sim **3.5**min

All 10000 instances on 32-core cluster: **409.5**min


There were:

6091 with **>0** probability (shown in histogram)

3909 instances with **0** probability (removed)

The highest probability was **0.63**

While star topology yields **0.91**



Conclusion

- Preliminary experiments indicate that distributed SMC in UPPAAL scales very nicely.
- More work to identify impact of parameters for distributing individual SMC?
- UPPAAL 4.1.9 available
(support for SMC, DSMC, 64-bit,..)

End